AIAA-98-4736

THE CHALLENGE AND PROMISE OF BLENDED-WING-BODY OPTIMIZATION

Sean Wakayama*
The Boeing Company
Long Beach, CA 90807

Ilan Kroo†
Stanford University
Stanford, CA 94305

Abstract

Multidisciplinary design optimization (MDO) is an important part of the Blended-Wing-Body (BWB) aircraft design process. It is a promising technology, but faces many challenges in routine application to aircraft advanced design. This paper describes current approaches, recent results, and future challenges for MDO as reflected in our experience with BWB design over the past four years. Current efforts have employed the Wing Multidisciplinary Optimization Design (WingMOD) code, targeting broad optimizations with large sets of design variables and constraints. These efforts have shown substantial payoffs stemming from the natural ability of MDO to handle the geometric complexity and the integrated design philosophy of the BWB. Challenges to MDO have been identified in the breadth and depth of the analysis desired to capture aerodynamic, stability, and control issues for this configuration. Future efforts include incorporating higher-fidelity codes while maintaining the breadth of scope, possibly with methods such as response surfaces and collaborative optimization.

Introduction

The Blended-Wing-Body (BWB) is a revolutionary concept for commercial aircraft[1-2]. It requires a design approach that departs from the conventional decomposition of the airplane into distinct pieces and instead integrates wing, fuselage, engines, and tail to achieve a substantial improvement in performance. This provides an arena rich in opportunities for multidisicplinary design optimization (MDO). The high level of integration breaks the normal design process; instead of satisfying specific requirements with a distinct airframe part, an array of requirements must be satisfied with an integrated airframe. This changes the design philosophy and requires developing experience in the new way of thinking. MDO presents a solution for these new design challenges. This paper describes some of the early application of MDO in the development of Boeing's BWB concept, focusing on the aerodynamic and structural optimization of the blended-wing planform and highlighting the opportunities for an expanded role of MDO in continuing design work.

Current State: WingMOD

MDO in the BWB program has been undertaken using several codes. Early conceptual and cabin layout optimization was carried out at Stanford University using both gradient-based and genetic algorithms; however, most of the current MDO work has been done with the Boeing Company's Wing Multidisciplinary Optimization Design (WingMOD) code. This code was originally developed for conventional wing and tail design, but has been adapted for use on the BWB. While this paper focuses on the use of WingMOD for BWB design, the basic conclusions regarding MDO in aircraft advanced design are more generally applicable.

Basic WingMOD Analysis

As described in References 3 and 4, WingMOD optimizes aircraft wings and horizontal tails subject to a wide array of practical constraints. It performs wing planform, thickness, and twist optimization, with design variables including overall span plus chord, sweep, thickness, and twist at several stations along the span of the wing. It also optimizes skin thicknesses, fuel distribution, spar locations, and control surface deflections. During optimization, WingMOD enforces constraints on range, trim, structural design, maximum lift, stability, control power, and balance.

WingMOD handles structural design and maximum lift constraints at a higher fidelity level than the traditional conceptual design process. It also incorporates stability, control, and balance considerations directly in the aircraft optimization, where the traditional conceptual design process handles these constraints outside the sizing loop. By performing detailed optimization while attending wide-ranging constraints early in the design process, WingMOD identifies ways to trade and maximize interdisciplinary advantages, generating well-rounded configurations that are usually achieved at great cost with traditional design processes.

---

* Senior Engineer/Scientist, Member AIAA

† Associate Professor, Member AIAA

To provide this capability, WingMOD employs analyses that have higher fidelity than those for conceptual design, but are faster than those generally associated with preliminary design. The basic WingMOD method models an aircraft wing and tail with a simple vortex-lattice code and monocoque beam analysis, coupled to give static aeroelastic loads. The model is trimmed at several flight conditions to obtain load and induced drag data. Profile and compressibility drag are evaluated at stations across the span of the wing using empirical data with lift coefficients evaluated from the vortex lattice code. Structural weight is calculated from the maximum elastic loads encountered through a range of flight conditions, including maneuver, vertical gust, and lateral gust. The structure is sized based on bending strength and buckling stability considerations. Maximum lift is evaluated using a critical section method that declares the wing to be at its maximum useable lift when any section reaches its maximum lift coefficient, which is calculated from empirical data. For trim, section zero-lift pitching moment is modified for trailing-edge deflections using empirical relations.

WingMOD fits within an advanced design process as sketched in Figure 1. The process begins with configuration and cycles through the disciplines, ending with a sized baseline after performance analysis. From the baseline configuration, WingMOD generates an optimized design. The airplane is analyzed in more detail than in the process developing the baseline. This includes explicit modeling of control surface deflections for trim and explicit calculation of span loading for weight and drag assessment. The optimized design can be cycled through an optional computational fluid dynamics (CFD) analysis to verify the aerodynamic predictions in WingMOD and to generate a true outer mold line. For faster cycle time with lower fidelity, the CFD analysis could be skipped. Either way, the optimized design is passed from configuration through performance analysis to validate the weight and performance estimates.

Genie Optimization Framework

Optimization services for WingMOD are provided by the Genie framework. Genie, a GENeric Interface for Engineering, was originally developed at Stanford University as a shell for performing generic engineering optimization problems. The idea behind its development was to build a single interface that was powerful enough to be used for most engineering problems yet simple enough to be linked with any analysis code.

The version of Genie used in WingMOD was modified at Boeing under NASA contracts to handle the requirements of several aircraft design optimization tasks. Efforts were made to develop features, which the original software lacked, that were needed on various optimization projects. Since most problems for Genie at Boeing could be cast as a single, integrated analysis, little was done to make it an integration tool with distributed computing capability; however, there are no obstacles to developing that capability. In its present form, Genie enables easy linkage between the analysis and optimizer, allows automated data calculation, provides data output in useful formats, provides information to facilitate scaling design variables and constraints, provides a selection of optimizers, allows flexible definition of optimization problems, and allows for the development of graphical user interfaces.

Enabling easy integration of new analyses was important in getting Genie to be used on more than one project. Linking an analysis to Genie involves writing a trivial analysis interface and communicating design data through simple data interface commands. The analysis interface takes commands from the command interface or the optimizer and simply calls the analysis with no arguments. The data interface provides simple functions that the analysis uses to get and put information from and to the database. Since these are software subroutine calls, programming is required to link an analysis to Genie. This may seem less attractive than communicating through files; however, the programming is very simple and pays for itself in faster data transfer between analysis and framework. For an all new analysis, data interface calls can replace traditional input-output, saving programming time.

Genie had automated optimization and calculation capabilities early in its development. Optimizations could be set up and run as background jobs on Unix platforms using a simple command language. To allow better visualization of the design space, the command language was expanded to allow calculations or optimizations at points in multiple dimensions to map objectives and constraints through the design space. While of limited use for wing design, this feature is very important for airplane sizing applications.

A complementary development was the capability to output results in formats for special graphic programs that generate multi-dimensional sizing thumbprints. More important for wing design problems, output capabilities were added to generate data summaries that could be rapidly inserted in spreadsheet programs to create detailed graphical reports that illuminate dozens of characteristics across the wing span.

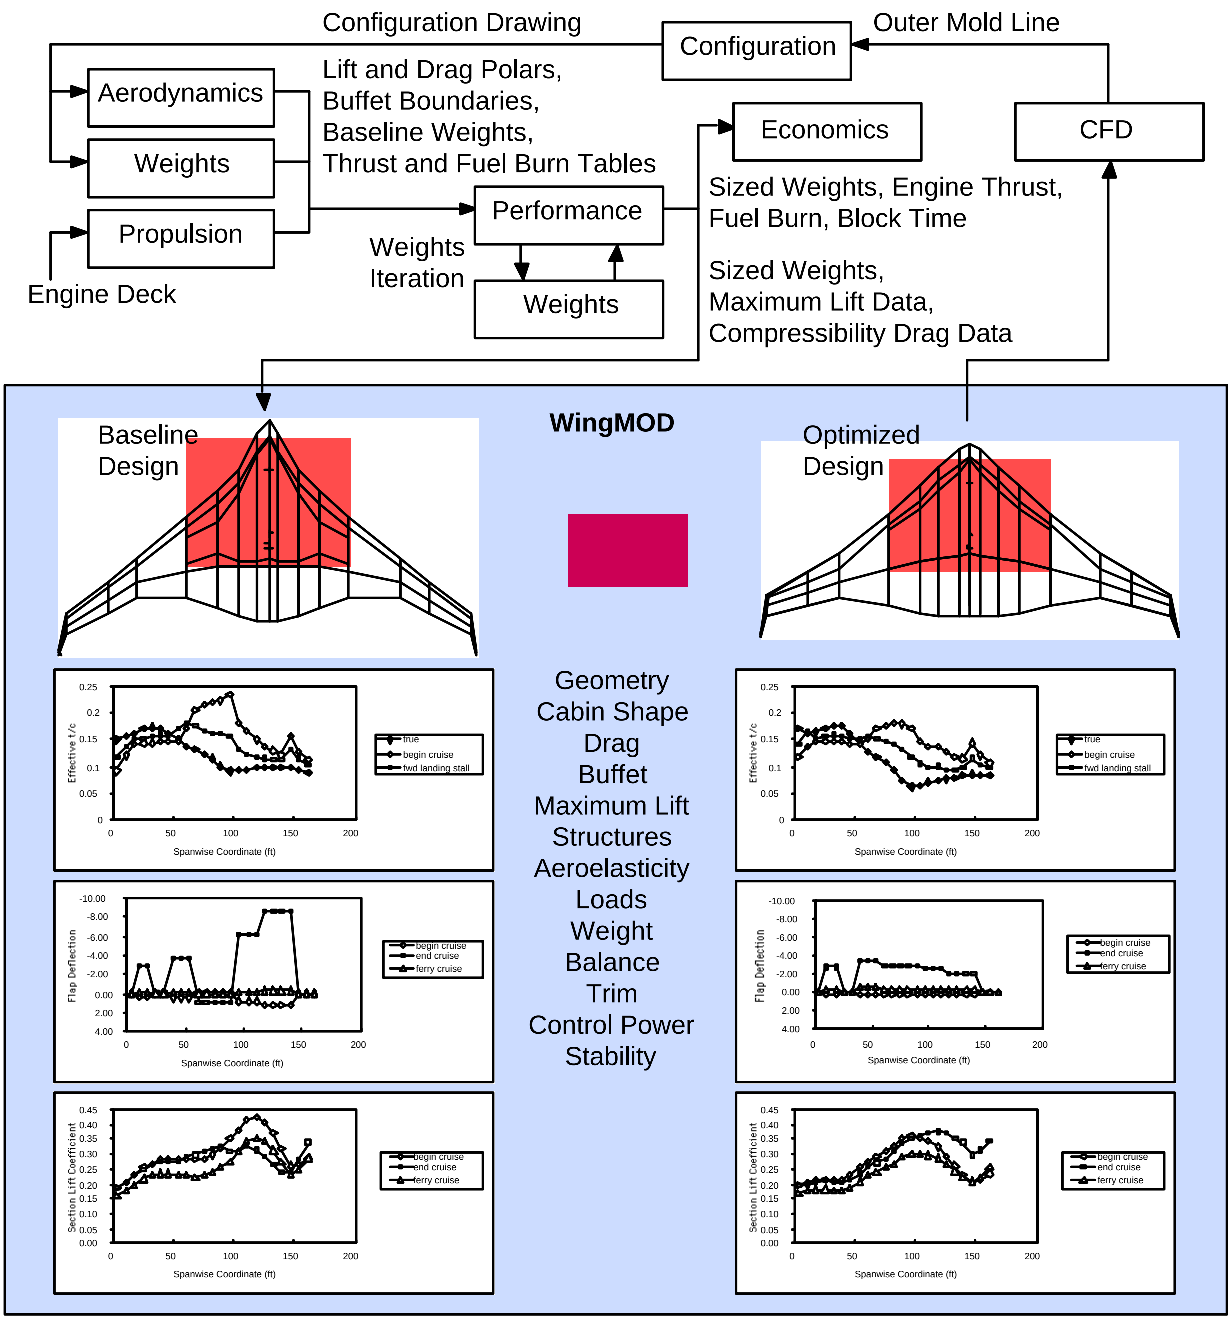

Figure 1. WingMOD design process

For detailed wing design problems, design variable and constraint scaling is extremely important for achieving timely, converged optimizations. Poor scaling can slow down or prevent convergence. The difficulty in selecting proper scaling comes from having a mix of very different variables and constraints that relate to each other in often non-intuitive ways. Very little is said about how to determine proper scales for all the variables in an optimization problem, and too often proper scaling is the result of a lot of experience by trial and error. There is a systematic approach to design variable scaling[4], which Genie facilitates through the Non-Linear Optimizer (NLOpt). NLOpt is based on sequential-quadratic programming and was written for use with Genie. At the end of each optimization, NLOpt provides information that can be used to improve scaling for subsequent optimizations. This feature has been essential to enabling optimizations in over one hundred design variables.

A motivation of using an optimization framework is the opportunity to make several optimizers available to the analysis. Genie provides access to the efficient NPSOL optimizer as well as the robust NLOpt. The switching between optimizers occurs within Genie where the analysis programmer does not need to worry about it. The optimizer is connected to a goal function interface, which acts like an ordinary function with arguments to the optimizer; however, the goal function interface works with the command and data interfaces to transcribe the abstract optimizer variables into physical variables. This way, the command interface can set any database variable to be a design variable, objective, or constraint, providing great flexibility in setting up optimization problems. The complex programming to provide this capability is invested in the framework while connections between analysis and framework are kept simple. This offers a large payoff for the low cost of linking an analysis with the framework.

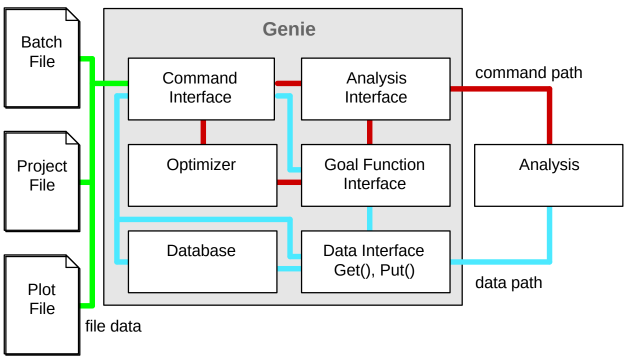

Figure 2. Genie optimization framework.

Graphical user interfaces (GUI's) provide a similar motivation for using optimization frameworks. While optimizations are run as Unix command line background jobs, Genie does have Macintosh and X-Window GUI's, which overlay the command interface. Investing in a GUI for a framework like Genie is attractive because that benefits every analysis connected to the framework. The challenge is then to create a generic GUI that can perform as well as application-specific GUI's for a range of analyses.

The combination of optimization framework and wing analysis make the WingMOD code. As described to here, WingMOD had been applied to design of a composite wing for a stretched MD-90[5] and for studies on the MD-XX. Application to the BWB would require substantial changes.

Challenges of the BWB

Radically different from conventional aircraft configurations, the BWB presents special design challenges. The integrated nature of the configuration is one challenge for which MDO offers a promising solution. Where the design of conventional aircraft can be divided between different disciplines, no discipline can work independently on the BWB. Where configuration can set the fuselage and aerodynamics can set the wing on a conventional aircraft, the two disciplines are forced to work together in defining a low-drag wing that adequately encloses the payload on the BWB. In that task, the large number of geometric degrees of freedom coupled with a number of geometric and aerodynamic considerations present a substantial MDO problem. Adding consideration of weight, balance, stability, and control issues turns this into an MDO challenge.

Further increasing the challenge, the BWB has unique design features that require higher fidelity modeling than might be acceptable for conventional designs. To enclose the payload within the wing, the BWB has very thick airfoil sections over its body. Attaining low drag, transonically, with these airfoils is an aerodynamic challenge. In this region, the wing structure doubles as pressure vessel for the cabin, presenting flat panels that must support pressure loads over large spans dictated by the cabin arrangement. Designing and analyzing these panels and assessing a weight for them is a substantial challenge for structures and weights disciplines. To reduce drag, the design is tail-less, but this creates interesting challenges for stability and control: first, to balance the airplane and provide sufficient control power, and second, to ensure that control deflections for trim do not adversely affect the spanload and hence the drag. A final challenge lies in the aft-mounted engines and the difficulties with propulsion and airframe integration. Before undertaking a credible MDO effort on the BWB, some of these issues had to be addressed with new analysis methods.

Aerodynamic Method Improvements

In aerodynamics, access to rapid Navier Stokes solutions has provided tremendous insight and confidence in the aerodynamic understanding of the BWB. The turn-around time for these solutions has been adequate for wing design in the cruise condition, allowing substantial progress in the aerodynamic design of the BWB. Unfortunately, these methods are not directly used by WingMOD. To touch on disciplines such as loads, low-speed aerodynamics, stability and control, WingMOD evaluates 20 flight conditions in each analysis. To explore a broad range of design changes, optimizations include over 100 design variables. With 20 flight conditions per analysis, 100 analyses per gradient calculation, and a minimal 100 major iterations of the optimization, we end up with 200,000 aerodynamic calculations per optimization. This strongly discourages any attempt to include a high-fidelity aerodynamic analysis directly within WingMOD.

The difficulty that severely delayed credible application of WingMOD on the BWB was the finding that the original WingMOD aerodynamics model was missing important characteristics that were captured in Navier Stokes codes. Because the flow around the center-body is three-dimensional, the center-body pressures correspond to the flow over a thinner effective 2-D section. 3-D relief is felt because the neighboring airfoils around the center-body are not as thick. This allows the thick sections that are needed to enclose the payload. In return, the outboard wing feels increased velocities because of the thick center-body and the pressures on its airfoils correspond to effectively thicker sections. These effects were modeled as described in Reference 6. An example of this effect is shown in Figure 3. Without this model, WingMOD could not produce aerodynamically feasible designs.

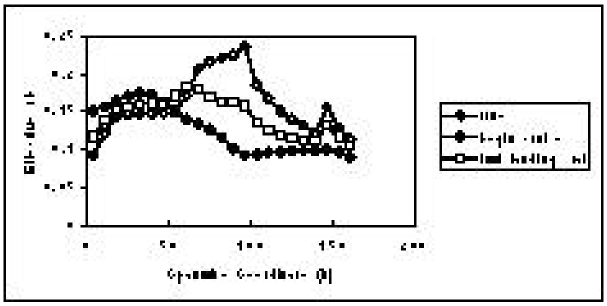

Figure 3. Baseline effective t/c distributions.

This example highlights a few of the obstacles to use of MDO in industry. First there is the reluctance to back-off on fidelity. Second is the breadth of criteria that should be considered in developing an optimal design. Coupled with the first obstacle, this either leads to prohibitively long optimization times or a substantial reduction in scope of the optimization problem. Third is a lack of intermediate-fidelity codes that can adequately substitute for high-fidelity codes at a fraction of the computing cost. The WingMOD approach tackles this third obstacle but continues to meet resistance on the issue of fidelity.

Structural Method Improvements

In structures and weights, a new method was employed to model the BWB center-body. The center-body is essentially wing structure, but it is pressurized and has very large rib spacing to accommodate the cabin. Structural equations were introduced in WingMOD to analyze wing skin panels as beam-columns with applied lateral pressure loads. This differs from the basic WingMOD buckling analysis, which looks only at buckling stability. Lateral pressure loads and compressive column loads from global bending moments are applied to the skin panels, generating nonlinear loads. Skin panels are modeled as sandwich structure with composite face sheets. While the core depth is set externally by manufacturability or damage tolerance constraints, the face sheet thicknesses are sized directly in the optimization to meet stress allowables. Panel stresses are evaluated at design running loads that are set in the optimization and are constrained to exceed actual running loads calculated through a wide array of structural design conditions.

Stability and Control Improvements

In the area of stability and control, the BWB forced the inclusion of new concerns in the WingMOD optimization, including scheduling control surface deflections and observing center-of-gravity issues. Scheduling control surface deflections is important because the airplane is trimmed with control surfaces distributed along the wing, which will impact the spanload and have first-order impacts on drag and weight. Center-of-gravity (CG) and balance issues are important because they indirectly affect the spanload by defining the trim points for the airplane.

To enable optimization of control surface deflections while emulating a realizable control law structure, WingMOD was modified to accept five deflection schedules: high-speed trim, high-speed control, low-speed trim, low-speed control, and maneuver load alleviation. These gear the control surfaces of all elements in the WingMOD model to pilot trim control, pilot maneuver control, and load factor. During optimization, control settings are set to trim the airplane and control surface gearing is selected to optimize performance. The high-speed trim gearing targets minimum trimmed cruise drag. The high-speed control and maneuver load alleviation gearings seek reduced critical loads. The low-speed gearings provide control authority over a range of conditions while preventing control surfaces from saturating or wing sections from stalling.

To assess center-of-gravity issues, WingMOD was modified to track the longitudinal position of structure, fuel, payload, and general discrete masses. The array of conditions analyzed in WingMOD includes conditions that set both forward and aft CG limits. During planform optimization, the limits are matched to the actual longitudinal balance. The range for the performance cruise mission is based on trimmed drag evaluated at the calculated CG. This encourages planforms that minimize CG range.

Propulsion Airframe Integration

Propulsion-airframe integration is an intimidating challenge for the BWB. This has been attacked through CFD analysis and inverse design, with initial results showing promise for solving the design problem albeit through a lengthy process. With the fine detail required for this work, there is little hope of incorporating this directly in a WingMOD optimization, although new approaches to course-grained distributed design are being investigated to accomplish this kind of integrated design capability[2].

Designing with WingMOD

With this brief description of the fundamental methodology of WingMOD and the design challenges of the BWB, we next summarize an example of the work that has been accomplished with WingMOD on the BWB. This example will hint at the detail and complexity that is needed to address an industrial aircraft design problem with MDO. This example shows the substantial gains that might be achieved on novel concepts, such as the BWB, where tight design integration and lack of design experience make the application of MDO not just a nicety, but a necessity.

Critics may argue that the problem addressed in this example is not broad enough or that the analysis methods are not deep enough to satisfy the concerns of industry. More is definitely desired in both breadth and depth, and much work remains to be done to achieve these improvements; however, WingMOD optimizations are providing answers that are useful to industry now. While the BWB program has yet to study an MDO-based design in detail, the directions taken by WingMOD in seeking optimal designs have provoked thought, discussions, and conventional studies that have led to improved designs. MDO has gained acceptance in the BWB program as a tool to find ways to improve the design.

This example uses a notional BWB developed under Task 18 of the Advanced Subsonic Technology (AST) program. The baseline airplane was configured and sized conventionally. The airplane mission was to carry 855 passengers 7,500 nmi at Mach 0.85, although less ambitious BWB configurations are currently under study. Further details of the optimization are given in Reference 6.

Design Conditions

To touch on most of the critical issues affecting the BWB, 20 design conditions were examined, as described in Reference 6. The BWB is highly sensitive to CG location because that governs the deflection of the control surfaces, the spanload, and ultimately the drag and weight. Where we can usually identify a critical CG location for each condition on a conventional airplane, the influence of control surface deflections on the spanload makes this difficult or impossible on the BWB; hence, several conditions are examined at both CG locations. This is one way the BWB stretches the breadth of any MDO effort. Even with this breadth, more conditions are desirable, with the first additions likely to be used for analyzing yaw control constraints.

Design Variables

Design variables are listed in Table 1. The details are discussed in Reference 6. The design variables cover both external geometry and interior arrangement of the major structural components. The boundaries of the cabin can be optimized as well as the distribution of fuel. Schedules for deflection of control surfaces and structural sizing can be handled. Optimizing these quantities results in a 134 variable problem.

This number of variables is admittedly small relative to some optimization problems (e.g. detailed structural sizing and trajectories through collocation); however, the extent of the geometric degrees-of-freedom make this an ambitious MDO problem. One obstacle to the use of high-fidelity codes in MDO has been the ability to automatically handle major geometry changes. FEM models and, to a lesser extent, CFD models would offer resistance to the planform changes examined in this example. The simpler models in WingMOD allow very broad variations in geometry to be explored. This is important for the BWB because there is too little experience with the design to substantially narrow the design space.

To those unfamiliar with the use of formal optimization in aircraft conceptual design, 134 design variables is quite a lot. It is more than a human would be able to sort out using conventional trade studies and exceeds the capability of most current advanced design codes.

Name	Number
mission takeoff weight	2
chord	9
sweep	7
t/c	8
incidence	7
payload chordwise extent	10
spar location	7
fuel distribution	6
nose tank fuel	3
CG limits	2
CG location	3
trim deflection schedule	8
control deflection schedule	8
trim angle of attack	16
trim or control deflection	16
trim load factors	2
design running load	13
center-body skin thickness	7
total	134

Table 1. Design Variables

This is particularly important for the BWB because existing tools that size thrust and wing area do not properly handle geometric changes to the BWB or account for important BWB constraints.

Constraints

Constraints are listed in Table 2. The constraints cover performance, stability, control, balance, structural design, buffet, and maximum lift. They also include geometric constraints that force the wing to wrap around a fixed payload. The details are left to Reference 6, but this table should indicate the breadth of constraints that are necessary to undertake an industrial MDO problem. There are a large number of constraints, 705, but only 90 are active. The constraint-based sequential quadratic programming algorithm used in WingMOD handles large numbers of constraints very easily, so the approach taken is to include all the constraints that could possibly drive the design and to let the optimizer determine the ones that do. When the active constraints are compared against the 134 design variables, there are 44 unconstrained degrees of freedom. This is a large dimension to explore that would take a prohibitively long time to navigate with conventional advanced-design methods.

Optimization Results

When the optimization was carried out, the design moved from the baseline configuration sketched in Figure 4 to the optimized configuration sketched in Figure 5. Additional human design input may be used to simplify the design from the optimizer, smoothing features that add much design complexity for small performance gains. This leads to a final design such as that shown in Figure 1. Alternately, the design could be re-optimized with fewer design variables after initial optimizations reveal the most important planform breaks.

The most substantial design changes were tighter packaging of the payload and the thinning of airfoils in the kink of the wing. By changing planform and thickness, the optimized design achieved a much tighter fit of the payload between the spars. The payload extent is indicated by the shading in the figures. This reduced the area of pressurized skin for a substantial reduction in weight. Thinning of the kink airfoil sections relieved compressibility drag penalties and allowed the optimized design to load the kink region for a better spanload and lower drag. This is described in more detail in Reference 6.

The final performance results are shown in Table 3.

Name	Number	Number Critical
range	2	2
CG location	3	0
CG limits	5	4

Table 2. Constraints.

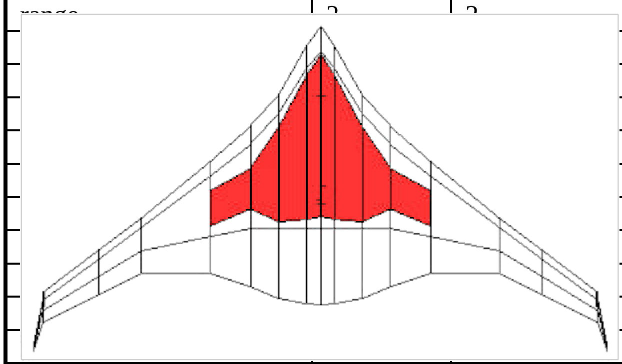

Figure 4. Baseline Configuration.

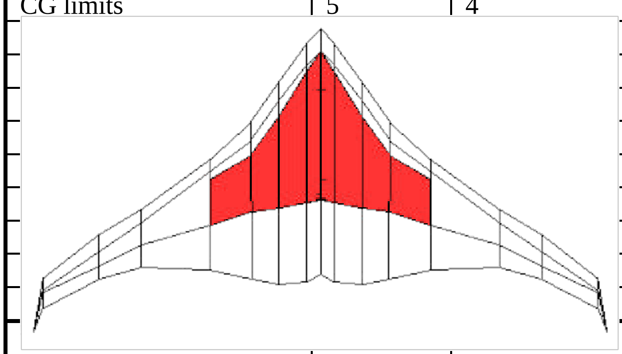

Figure 5. Unmodified Optimized Configuration.

Operating empty weight is reduced by better packaging of the payload. L/D is increased, largely because of better span loadings. The baseline airplane has a WingMOD-optimized span load that balances weight, drag, and control considerations. Had the baseline span load been aerodynamics-optimized, the optimized design would show little improvement or even degradation in L/D, but it would show more substantial empty weight reduction. In optimizing from aerodynamics-defined wings, MDO almost always finds ways to improve the other disciplines at a small expense to aerodynamics. This can make it difficult for MDO-based designs to gain acceptance from aerodynamics, especially because the aerodynamic penalties can be captured with high-fidelity early in the design process while the projected gains in other disciplines can take months to substantiate.

performance figure	% change from baseline
takeoff weight	-6.9
operating empty weight	-5.0
fuel burn	-12.0
gross area	+0.8
average L/D	+7.5

Table 3. Optimization Results

The combination of weight and drag reduction results in substantial reductions in fuel burn and takeoff weight. Wing area, which is a primary design variable for conventional sizing methods, is virtually unchanged, meaning that improvements were made through much finer manipulation of the geometry. This shows a fundamental advantage of multidisciplinary optimization over conventional sizing processes. In addition, the design was accomplished in a short time, with overnight optimization runs and a few tries to perfect the optimization problem. This contrasts with the months of study that would be required to optimize the design conventionally. The design improvements and speed that MDO offers show great promise for advancing BWB design.

The Promise of MDO

The basic conclusion of this exercise is that the design capabilities of an MDO process can lead to substantial improvement in the design of a novel configuration such as the BWB. There are some less-visible advantages that come from MDO codes that are described below.

Design Cycle Time

In design studies using conventional methods, the following observations could be made. The conventional advanced design process uses 3 to 6 weeks for a BWB planform change to cycle through configuration, weights, aerodynamics, and performance analysis. Optimizing an aircraft could take several cycles (months) to optimize the aircraft. The cycle time limits the number of design variations that can be explored. Even worse, this cycle time only covers performance analysis; additional time is required for balance, aeroelastics, stability and control.

In the course of this study, the advantages of the WingMOD approach could be seen. It still takes 3 to 6 weeks to model and calibrate a baseline design in WingMOD. This is comparable to the cycle time for a planform analysis using conventional methods; however, only a single run is needed to optimize the aircraft, reducing months of cycle time to an overnight job. In addition, the optimization handles many more design variations than could be explored by the conventional methods. Finally, the optimization deals with balance, aeroelastics, stability and control issues that the conventional approach leave for later analysis and revision.

First-Cut Information

To perform multidisciplinary optimization, coupling aerodynamic loads and structural design is almost a must. From there, it is natural to make that an aeroelastic calculation. Doing this for MDO adds the advantage that the resulting system is highly-automated, fast, and robust. An unexpected benefit is that an MDO code, like WingMOD, with grandiose expectations of planform optimization becomes amazingly useful for mundane tasks such as providing a first estimate of loads, an initial sizing for skin thicknesses, and aeroelastic stability data. While there are industrial processes in place to do all this, they are expensive, time-consuming, and have a chicken-or-the-egg problem: how do you generate loads when you need skin thicknesses to capture the aeroelastic effects, but you need loads to figure out what the skin thicknesses need to be? The fidelity of those processes justifies their expense, and we would never use WingMOD to certify an airplane, but WingMOD is perfect for getting the first cut at the loads and structural sizing from which the detailed processes can start.

Individual Versus Total Good

Pushing for overall airplane improvement over individual discipline improvement can be a difficult practice to incorporate in a large design team, but it is especially important for a revolutionary concept such as the BWB. Traditionally, aerodynamics has taken the lead in defining wing shape. A compelling reason for this is the speed of aerodynamic processes: a wing-only planform change can be put through CFD in as little as a few days. The other disciplines are not so lucky: a finite element model can take six months. So while aerodynamics can push for a particular design with hard facts, the other disciplines can offer only qualitative objections, or the design cycle must drag on for hard numbers from the other disciplines. For conventional aircraft, this is not very critical: the wing turns out heavier and the tail turns out bigger than they ought to be, but the airplane will still work. For the BWB, this could be disastrous: the wing shape that maximizes L/D is unlikely to lead to a balanced airplane with the control authority to rotate for takeoff. The WingMOD approach looks at all the design drivers it can to offer a design that is the best compromise between the disciplines. Analyses provide hard numbers, albeit approximate, for each discipline, making it difficult for any one discipline to dominate. It is difficult to accept that WingMOD designs inevitably come in with lower L/D than aerodynamics group knows they can achieve, while offering benefits in other disciplines that cannot be immediately verified. Even within aerodynamics, WingMOD will compromise cruise performance to enable meeting low-speed lift and control requirements.

Future Directions

While WingMOD optimization has made promising first steps toward solving the BWB design problem through MDO, much more is desired. Problem areas specific to the BWB are identified below.

Increased Breadth

While the breadth of conditions examined by WingMOD is relatively well accepted, there are instances where more is desired. An example is modeling engine-out lateral control. From experience with the BWB-17 Flight Control Testbed[2], this could drive the sizing of the outboard wing and winglet chords, which affect the effectiveness of the rudders needed to control this condition.

Higher-Fidelity Codes

The WingMOD aerodynamics module certainly leaves something to be desired for analyzing the BWB; however, the speed of this analysis is required to cover the breadth of flight conditions that are essential to performing any multidisciplinary planform optimization. Because of their speed, higher-fidelity panel methods are the most likely next-step to improving the WingMOD aerodynamic analysis. Incorporating a true CFD analysis promises the benefit of capturing all the important aerodynamic effects and the ability to directly handle the propulsion-airframe integration problem; however, direct inclusion of CFD at this time would likely bring a WingMOD-breadth optimization to a screeching halt.

The inclusion of finite element methods (FEM) is a lower priority than CFD. This is because the span time for generating adequately detailed FEM models is too long for them to be used actively in the conventional design process. Design work on the BWB uses weight estimates from parametric equations that may be calibrated to but are really independent from FEM results. The intermediate-fidelity structural analysis in WingMOD is already better than parametric weight equations, so the optimization cannot be faulted with missing something the standard approach would catch. At this stage, FEM work is very important for calibrating weights codes and verifying that there are no show-stoppers in the design, but it works too slowly to substantially impact planform trade studies. If FEM analysis had a span time equivalent to that for CFD analysis, then it would play a stronger role in the early definition of an aircraft, and there would be a greater impetus to include it in advanced-design MDO.

A tantalizing prospect for increasing the fidelity of a WingMOD-type optimization is incorporation of a detailed mission analysis code. This could bring high quality to the performance figures at little computational cost. It could also eliminate many stand-in constraints, for example takeoff speed targets instead of a true field length constraint. The issues here include judiciously selecting a minimum number of aerodynamic analyses to provide the data required for the mission analysis and generating noise-free numbers from the mission analysis.

Propulsion-airframe integration is especially important for the BWB because of the potential for either improved performance or problems with high distortion associated with boundary layer ingestion. Because of the complex, viscous, transonic flow in this region, simple models are ineffective and one is forced to rely on rather time-consuming CFD simulations for reliable guidance. The simple framework on which WingMOD is based is not well-suited to the incorporation of such methods and future work is clearly required in this area.

Optimization Framework Improvements

While the example optimization presented in this paper is large, many other parameters must be input to run WingMOD and this presents an often bewilderingly steep learning curve. Improvements in the way the framework handles large numbers of variables would help divide the problem to be more tractable to the user.

Applying techniques for decomposition through the optimization framework would be ideal. That would provide additional capability while allowing sub-problem analyses to remain unchanged and unburdened by the complexity of the overall optimization problem. A candidate project likely to help BWB optimization studies would build collaborative optimization capability into the Genie framework.

The Challenge of MDO

The promise of MDO has been suggested by our recent experience with BWB design; however, it has also highlighted some of the generic problems and challenges in industrial acceptance of MDO.

Problem Formulation

As with single discipline optimization, correct and efficient problem formulation is critical to obtaining useful results from MDO. Because of the subtle interactions and interdisciplinary feedback that may be less well known to disciplinary experts, it is often more difficult to anticipate the weaknesses of analyses or the ill-posedness of a particular problem. Experience with both conventional and BWB design suggests that the problem formulation (selection of design variables, objectives, constraints, and bounds) evolves as the design is developed. It is naïve to expect that a realistic large-scale MDO problem can be fully-understood by a design team from the outset. While automatic aerodynamic optimization with a specified planform and a restricted set of design conditions can be reasonably well formulated a priori, the multidisciplinary aircraft design problem is more challenging and calls for a qualitatively different approach. One must structure the problem in such a way that changes in design variables and constraints can be made along the way. Individuals and truly integrated teams must routinely meet to evaluate the results and refine the analysis requirements or problem definition. The potential for impractical designs that reduce the credibility of the process is great without such planned intervention.

Breadth Versus Depth

Multidisciplinary optimization results are often criticized for being so limited in scope or fidelity as to be merely academic exercises—and such criticism is often well justified.

Based on the notion that a chain is only as strong as its weakest link, low-fidelity models covering many disciplines are sometimes omitted, leaving a two or (rarely) three discipline MDO problem that uses sophisticated disciplinary models. A chain with missing links is worse than one with weak links. A classic example is that aerodynamic and structural optimization without consideration of maximum lift leads to wings with absurdly small tip chords[3]. The BWB design problem illustrates the large number of disciplines to yield reasonable results.

On the other hand, the BWB represents an example of a design for which 2-D section analysis superimposed on a simple 3-D model fails to reveal some of the fundamental opportunities available in the BWB design space. The use of too-simple analyses might lead one to conclude that the advantages of the concept were insufficient to warrant the development of improved analyses or further consideration.

This is one of the most fundamental dilemmas in MDO that will not be solved by advances in optimization theory or AI. Practical MDO will always require good engineering judgement to match the scope of the particular problem to appropriate analyses. Approximate models are often very adequate, depending on the actual sensitivities of active constraints and objectives to the particular choices for design variables. Rapidly increasing computational capabilities including parallel systems and efficient algorithms will change the selection of appropriate models, but will not reduce the importance of this step. As more sophisticated analyses become feasible, the importance and difficulty of problem formulation and integration will only increase.

Optimization Analysis Requirements

The breadth versus depth problem would be alleviated if high-fidelity analyses ran faster. Intriguing options for increasing CFD optimization speed are automatic differentiation and adjoint formulations, the latter promising sensitivity information for little more than a function evaluation, although the present problem involves a large number of constraints that reduce the attractiveness of an adjoint approach. The best improvements for FEM lie in automating the model generation process.

Beyond speed, analyses must be robust and smooth to be used with optimization. The robustness of automatic grid generation through large planform variations is a problem. The smoothness of CFD and FEM results is also an issue.

Speed, robustness, and smoothness are also an issue with mission analysis codes. While several mission analysis codes exist that admirably fill the requirements of engineering analysis, the requirements for optimization motivate the creation of new codes that are built for optimization from the ground up. Such codes could use techniques, such as collocation, that make sense for optimization but were not important for the engineering needs the existing codes were written for.

Integration

In the development of WingMOD little attention was given to allowing for integration of existing codes or optimization decomposition techniques: the lack or unavailability of fast, intermediate-fidelity codes made it more expedient to develop an all-new, tightly-coupled analysis, which would not benefit from decomposition.

As more complex aerodynamic, structural, and dynamic analyses are included in BWB optimization, the basic tightly-integrated framework on which WingMOD is based begins to become unwieldy. Several research programs are currently underway to address such problems, although applications as complex as the BWB planform design problem have not been satisfactorily demonstrated to date. This would constitute an excellent test for the industrial applications of various concepts for decomposed analysis and distributed design, Reference 7.

The best near-term possibilities for bringing CFD into BWB MDO may be in the use of response surfaces and collaborative optimization. Collaborative optimization would isolate the CFD analysis in its own sub-space. Response surface techniques could map and capture the CFD analysis sub-space, which could then be included in a collaborative optimization formulation with WingMOD capturing the non-aerodynamic disciplines. Alternately, response surfaces could simply capture aerodynamic data from specific CFD runs to be laid over WingMOD aerodynamic results. A tight coupling of aerodynamics and structures for aeroelastic loads calculation combined with a loosely-coupled, higher-fidelity aerodynamic performance code may solve some of the problems that involve both high dimensionality coupling and the need for very accurate aerodynamic solutions.

Although various techniques for loosely coupling multidisciplinary design problems have been proposed, (e.g. concurrent subspace optimization and collaborative optimization[8-9]), few have seen application in industry projects. We attribute this primarily to the fact that these techniques are still the subject of active research and have not matured to the point that they are easily implemented as an option in a commercial software package. The availability of such technology may reduce the need for an individual who understands both the particular design problem and the details of the optimization framework and theory. Although progress in this area continues, Reference 10, we do not expect that such a system is imminent.

Validation

There are very few examples of MDO-derived designs being validated to the point of being real, useable configurations. Advanced-design level optimization needs to be validated with high-fidelity analysis; high-fidelity optimization needs to be validated through broad analysis checks. Reference 5 describes the use of WingMOD to develop a conventional aircraft wing configuration and the subsequent CFD validation. Achieving acceptance for MDO in industry will require more examples of validated optimized designs. No validation of an optimized BWB design has been done, but WingMOD designs are close to being assessed with CFD. Passing the challenge of validation will be most important to bringing MDO to the forefront of BWB design.

Conclusions

The BWB is a revolutionary concept that benefits from MDO and yet illustrates the many challenges to its use in industry. Current efforts with the WingMOD code have been stretched in depth, particularly to capture unusual aerodynamic characteristics, and in breadth, to capture stability and control issues. Introducing high-fidelity analysis would be highly desirable, a prerequisite for handling propulsion-airframe integration, yet the breadth of the BWB design problem almost prohibits the direct substitution of more sophisticated codes for the current simpler models.

Much progress has been made with the advanced-design level WingMOD code. Successful optimization has been made with a large, comprehensive set of design variables and constraints. Attacking this broad problem has offered substantial payoffs because of the youth of the BWB concept: current BWB configurations are not as finely evolved as conventional transports. The success in handling this broad design problem has partly been facilitated through capabilities provided by the Genie optimization framework.

MDO offers much promise for improving the BWB. Optimization studies have shown potential for substantial reductions in takeoff weight. This comes from the ability of MDO to handle many more degrees of freedom and track more interactions across disciplines than conventional advanced-design processes. The BWB can benefit greatly from MDO because of the complexity of its geometry and the integrated nature of its design. In addition, the innate automation required for optimization offers significant reductions in design cycle time while handling considerations beyond the scope of the existing processes, including control surface deflections, balance, control, and aeroelastic effectiveness.

Achieving the promise will involve more work. Increased breadth of analysis and optimization framework improvements will evolve naturally, although it would be desirable to accelerate those developments. Incorporating higher-fidelity codes while maintaining the breadth of scope will be a large challenge, offering opportunities to test methods such as response surfaces and collaborative optimization. While current BWB work demonstrates the potential for MDO in aircraft advanced design, it remains to verify the predicted advantages of these optimized designs using more refined analysis codes.

References

[1] Liebeck, R. H., Page, M. A., Rawdon, B. K., "Blended-Wing-Body Subsonic Commercial Transport," AIAA Paper 98-0438, Jan. 1998.

[2] "Blended-Wing-Body Technology Study," Final Report, NASA Contract NAS1-20275, Boeing Report CRAD-9405-TR-3780, Oct. 1997.

[3] Wakayama, S., Kroo, I., "Subsonic Wing Planform Design Using Multidisciplinary Optimization," *Journal of Aircraft*, Vol. 32, No. 4, Jul.-Aug. 1995, pp.746-753.

[4] Wakayama, S., *Lifting Surface Design Using Multidisciplinary Optimization*, Ph.D. Thesis, Stanford University, Dec. 1994.

[5] Wakayama, S., Page, M., Liebeck, R., "Multidisciplinary Optimization on an Advanced Composite Wing," AIAA Paper 96-4003, Sep. 1996.

[6] Wakayama, S., "Multidisciplinary Design Optimization of the Blended-Wing-Body," AIAA Paper 98-4938, Sep. 1998.

[7] Kroo, I., "Multidisciplinary Optimization Applications in Aircraft Preliminary Design—Status and Directions," AIAA Paper 97-1408, Apr. 1997.

[8] Kroo, I., "Decomposition and Collaborative Optimization for Large-Scale Aerospace Design Programs," in *Multidisciplinary Design Optimization: State of the Art*, N. Alexandrov and M. Y. Hussaini, editors, SIAM, 1996.

[9] Braun, R. D., and Kroo, I. M., "Development and Application of the Collaborative Optimization Architecture in a Multidisciplinary Design Environment", in *Multidisciplinary Design Optimization: State of the Art*, N. Alexandrov and M. Y. Hussaini, editors, SIAM, 1996.

[10] Sobieski, I. P., Manning, V. M., Kroo, I. M., "Response Surface Estimation and Refinement in Collaborative Optimization," AIAA Paper 98-4753, Sep. 1998.